

SAM (REX/Ops) Operations Meeting 2008.07.15

Adam (typist), Eric, SteveW, SteveS, Jon Yoh, Angela, Randolph, Art, Robert, Anil

Meeting time: 45

Posting these minutes closes

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

No new releases.

Releases in the works (fixed means in CVS, but not released into Kits)

The screenshot shows two JIRA release notes sections. The first section is for 'sam_db_srv_pkg_v8_5_0' with a progress bar indicating '2 of 3 issues have been resolved'. It lists three issues: SAMDEV-8 (UNRESOLVED) about the 'samadmin remove cached files' command, SAMDEV-12 (FIXED) about 'Change D0 autodeestination logic', and SAMDEV-9 (FIXED) about an 'Incorrect exception specification'. The second section is for 'sam_v8_4_2' with a progress bar indicating '3 of 3 issues have been resolved'. It lists three issues: SAMDEV-11 (FIXED) about 'SAM client not properly threadsafe', SAMDEV-14 (FIXED) about 'sam generate recovery project', and SAMDEV-16 (FIXED) about 'Improve documentation'.

sam_db_srv_pkg_v8_5_0 (Release Notes) Progress: 2 of 3 issues have been resolved

- [SAMDEV-8](#) UNRESOLVED *samadmin remove cached files command does not copy the cached files record to the history table*
- [SAMDEV-12](#) FIXED *Change D0 autodeestination logic for test versions*
- [SAMDEV-9](#) FIXED *Incorrect exception specification for DataFileImpl.removeReplica()*

sam_v8_4_2 (Release Notes) Progress: 3 of 3 issues have been resolved

- [SAMDEV-11](#) FIXED *SAM client not properly threadsafe*
- [SAMDEV-14](#) FIXED *sam generate recovery project provides no way to only ask for the unconsumed files*
- [SAMDEV-16](#) FIXED *Improve documentation for "sam modify request"*

Issues not assigned to a release

T	Key	Status	Summary
	SAMDEV-5	Open	Improve project robustness against DB failures
	SAMDEV-1	Open	Convert station to use omniORB
	SAMDEV-17	Open	Broken check for modifying global.requestId in RequestImpl.updateRequestParams

DØ:

July 10: Shifter Mundal reported: "Noticed on the DHAtAGlance webpage that DB Servers not responding was -999 for a time. However, this number went to 0 after a while." --- Just mentioned it here (I've never seen this).

???? We don't know what this means either. Probably an error code meaning that he couldn't get any information at all. Networking or overloading problem with the machine?

July 14: Memory for the tomcat for the Dataset Definition Web interface has been increased to 1 GB (after the complaint of DZSAM-144 ---> pointing to problems when the no. of files was 21777).

Is this (probably very old) version of Tomcat leaking memory?

Possible plan - upgrade Tomcat or move it out of tomcat. Maybe make it some fancy Ajax thing (would have been good for a summer student).

Recent unresolved issues :

DZSAM-142 : "farm job submissions halted" -- M. Diesburg ---> the limit of forwarding nodes has been increased to deal with this --- apparently a major problem for Mike.

Actually I did not increase the number of slots - the number of running jobs dropped to an ok level.

DZSAM-141 : "MC jobs held again" affecting MC jobs in OSG Grid -- J. Snow : CPU spikes observed... no solution yet (as it seems to me).

Part of this is the periodic expressions "outage" but there's another daily outage too that we don't understand. Long term solution is another queueing node for Joel and going to Condor-7.

DZSAM-140 : Complaints about samgrid monitoring page --- samgrid.fnal.gov page seemed to be found. Resolved ??

Page came back - another condor outage??

Robert: d0srv058 is dead (one of the cab cache). Hardware fault and it won't boot -- service call placed. This is maybe 8 TB cache server. Robert told FEF about it.

Adam: Copy of Joel's OSG report...

DZero MC OSG production rose 2.0M events to 4.2M events last week. This is 14% above the average weekly OSG production rate for the past year of 3.7M events.

Production characterization: Good (with commentary).

Two VO specific issues negatively affected OSG production and efficiencies. First was many jobs failing because of missing or invalid parameter specifications for the generator application being run.

This quality control failure happens infrequently (but more than I'd like to see) and occurs upstream of production. This problem affected production for a couple of days at the beginning of the week. Second was related to samgrid infrastructure. Symptom is many grid jobs go to held status at forwarding nodes, bunched in time. This has been occurring repeatedly for several days.

Investigations suggest it is related to a "feature" of Condor 6 that is not a problem in Condor 7. This problem causes large inefficiencies in local job running and in other areas of production including

data production. From my point of view the Condor 6 -> 7 upgrade should occur ASAP. On the positive side D0 experienced good production results from OSG despite the VO specific problems.




Andrew has noted there are some problems with the dcaches at UNL and IU. A ticket will be opened and he will be in tomorrow to address them.

local job efficiencies:







About the same to a little better than last week. For a general OSG cluster I am seeing 3-10% job efficiency. Again optimistically speaking, in light of the VO specific problems mentioned above this result is not so bad. With regard to the effect of the new SE at MWT2-IU, the consolidated numbers per site obscure the story due to the VO problems mentioned above. Looking at the efficiencies of individual jobs at MWT2-IU, the story becomes clear. Many grid jobs at MWT2 had zero efficiency. This was do to bad requests running.









Looking at jobs with nonzero efficiency it can be seen that they are in general very high. I conclude that the addition of the SE has greatly improved efficiency at MWT2-IU. The strategy of adding SE's to significant production sites should be pursued with vigor.

Issues: DZSAM resolved issues from past week (Displaying 3 of 3)

	DZSAM-145	scheduler samgrid.fnal.gov is not found
	DZSAM-144	Re: error when using sam
	DZSAM-143	Unable to find specific event.

Issues: DZSAM new issues from past week (Displaying 6 of 6)

	DZSAM-145	scheduler samgrid.fnal.gov is not found
	DZSAM-144	Re: error when using sam
	DZSAM-143	Unable to find specific event.
	DZSAM-142	farm job submissions halted
	DZSAM-141	MC jobs held again
	DZSAM-140	samgrid is in hiding...

Issues: DZSAM unresolved problems (Displaying 8 of 8)		
	DZSAM-142	farm job submissions halted
	DZSAM-141	MC jobs held again
	DZSAM-140	samgrid is in hiding...
	DZSAM-129	MC production jobs unable to get files on cab2
	DZSAM-100	MC Jobs overwhelming SBGRID
	DZSAM-62	40 MC jobs held over the weekend
	DZSAM-54	Jobs don't work on Lancaster Forwarding Node
	DZSAM-53	samgrid reporting incorrect status of jobs

CDF:

Generally quiet shift, w/o major problems. Most notable events:

- Wed./Thu. July 9/10: dcache pools fcdldata146 and fcdldata147 offline due to memeory consumption problems (see CDFDCACHE-3 for more informations), Angela/Alex are working on it: open issue.

Notorious fcd146 is being tested by FEF. Angela says that FEF will try to replace the mother board even if they don't find problems with their tests. Plan for 147 depends on what happens with 146.

- Thu. July 10: CDFDH-171: tapes IAB128 and IAF532 in "NOACCESS" state, fixed by Stan: resolved issue.













- Fri. July 11: CDFDH-172: disk problem with b0csl23. Online experts are working on it: on going.

- Sat. July 12: mover 9940B22 finally back online (bad system disk): resolved issue.














- Sat. July 12: CDFDH-175: dcache pool fcdldata097 offline (not pingable), fixed by rebooting (Etta/Mark) next monday (help-desk ticket 118598): resolved issue.


















- Sat./Sun. July 12/13: CDFDH-176: Tapes IAB901 and IAD746 in "NOACCESS" state, then same happened for IAB812 and IAF572, message sent to enstore-admin: resolved issue.

Issues: **CDFDH resolved issues from past week** (Displaying 12 of 12)

	CDFDH-180	[JIRA] Account created
	CDFDH-177	SAM registration
	CDFDH-175	dcache pool fcdldata097 not pingable
	CDFDH-172	Warning: Critical Error on b0csl23.fnal.gov
	CDFDH-171	Tapes IAB128 and IAF532 in "NOACCESS" state
	CDFDH-170	Reading files
	CDFDH-169	HelpDesk ticket 118426 has been assigned to you NAKAMURA, KOJI
	CDFDH-168	Error on upload server fcdldata322 due to unregistered dataset uhgp6
	CDFDH-148	Added disk inactive from long time
	CDFDH-147	sam get dataset offsite
	CDFDH-142	dcache error for cafjob 3531028?
	CDFDH-109	SAM problems @ CNAF

Issues: **CDF new issues from past week** (Displaying 13 of 13)

	CDFDH-190	SAM Registration
	CDFDH-189	SAM registration
	CDFDH-180	[JIRA] Account created
	CDFDH-178	Fwd: Rejected posting to CDF_DISKPOOL_ADMIN@LISTSERV.FNAL
	CDFDH-177	SAM registration
	CDFDH-176	4 CDF tapes in "NOACCESS" state
	CDFDH-175	dcache pool fcdldata097 not pingable
	CDFDH-173	reading problem for cdfdca1.fnal.gov data from caf
	CDFDH-172	Warning: Critical Error on b0csl23.fnal.gov
	CDFDH-171	Tapes IAB128 and IAF532 in "NOACCESS" state
	CDFDH-170	Reading files
	CDFDH-169	HelpDesk ticket 118426 has been assigned to you NAKAMURA, KOJI
	CDFDH-168	Error on upload server fcdldata322 due to unregistered dataset uhgp6

Issues: CDFDH unresolved problems (Displaying 17 of 17)		
	CDFDH-190	SAM Registration
	CDFDH-189	SAM registration
	CDFDH-178	Fwd: Rejected posting to CDF_DISKPOOL_ADMIN@LISTSERV.FNAL
	CDFDH-176	4 CDF tapes in "NOACCESS" state
	CDFDH-174	hi, I am trying to submit a skimming job but...
	CDFDH-173	reading problem for cdfdca1.fnl.gov data from caf
	CDFDH-165	Problems with file transfers - please investigate
	CDFDH-164	/pnfs/cdfen mount on ILP
	CDFDH-162	DH performance ?
	CDFDH-150	No authroization request
	CDFDH-139	Long running query from SAM on CDF Offline Production database
	CDFDH-138	Problems with fcdpdata147
	CDFDH-134	sam dataset definition problem
	CDFDH-128	cdf-caf and cdf-fncdfrsv0 stations can get stuck if restarted
	CDFDH-121	Debugging SAM @ CNAF
	CDFDH-112	re : fcdpdata135 is offline
	CDFDH-96	cdf-japan station is down

MINOS:

General

Beam returned on Sunday 13 July, down since early June.

Enstore -

DCache -

SAM -

Oracle - Scheduled a 'taking stock' meeting at 09:00 today.

OPEN ISSUES

IT 1146: cannot test existence of sam tape location
IT 2071: quadratic project startup time growth

Art wants to know about SL5. Robert has run the stager component on SL5 and it works. CDF is running on SL5.0 with the nameservice, station and DB servers. Robert did discover an issue - service principals would not authenticate. Fix was for Joe Boyd to remove some stuff from the config file.

PLAN -

DB:

D0 :

Alert(S):

- Information for: d0ofprd1

Increased the max extents to 20, 000 for SAMPRD PROJECT_FILES TABLE
9996

- Information for: d0ofprd1

Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
-------	--------	------	--------------	-------------------	--------------

SAMPRD	ANALYSIS_PROCESSES	TABLE	596	600	4
--------	--------------------	-------	-----	-----	---

Action: Increased the max extents to 1200

CDF :

Alert(s):

Angela pulled back on the cron job for integration monitoring plots.

MINOS:

Alert(s):

No date set yet for migrating to Oracle 11?
Also have to worry about kerberos migration.

Enstore/dCache

Eric: 2 LTO3 tape drives were added for CDF (thanks to Stan!).

Angela has put the dCache JIRA project on the back burner - will use helpdesk instead. Sent two main tickets to the helpdesk. Not much has been done.

Timur complained about JIRA spam to the helpdesk and created his own spam-storm for other people. Unfortunately a user was added with the e-mail of a ticket generating e-mail list. Angela removed the user and stopped the looping.